

# Evolving 5G: ANIARA, an Edge-Cloud perspective

Ian Marsh  
RISE AB, Sweden

Wolfgang John  
Ericsson AB, Sweden

Ali Balador  
Ericsson AB, Sweden

Federico Tonini  
Chalmers University, Sweden

Jalil Taghia  
Ericsson AB, Sweden

Andreas Johnsson  
Ericsson AB, Sweden

Paolo Monti  
Chalmers University, Sweden

Jonas Gustafsson  
RISE AB, Sweden

Pontus Sköldström  
Qamcom AB, Sweden

Johan Sjöberg  
Ericsson AB, Sweden

Jim Dowling  
Hopsworks AB, Sweden

## ABSTRACT

Emerging use-cases like smart manufacturing and smart cities pose challenges in terms of latency, which cannot be satisfied by traditional centralized networks. Edge networks, which bring computational capacity closer to the users/clients, are a promising solution for supporting these critical low latency services. Different from traditional centralized networks, the edge is distributed by nature and is usually equipped with limited connectivity and compute capacity. This creates a complex network to handle, subject to failures of different natures, that requires novel solutions to work in practice. To reduce complexity, more lightweight solutions are needed for containerization as well as smart monitoring strategies with reduced overhead. Orchestration strategies should provide reliable resource slicing with limited resources, and intelligent scaling while preserving data privacy in a distributed fashion. Power management is also critical, as providing and managing a large amount of power at the edge is unprecedented.

## CCS CONCEPTS

• **Hardware** → **Energy metering**; • **Computing methodologies** → **Multi-agent systems**; **Self-organization**.

## KEYWORDS

Edge computing, federated learning, container technologies and orchestration, energy metering.

## 1 INTRODUCTION

The project is based on two use case families addressing smart manufacturing and smart cities. Figure 1 illustrates the functional technologies for evolving 5G edge systems. The application scenarios are drawn upon the synergies between cloud computing, 5G, and mobile edge compute.

Pontus Sköldström was with Ericsson AB at the time of writing this paper.

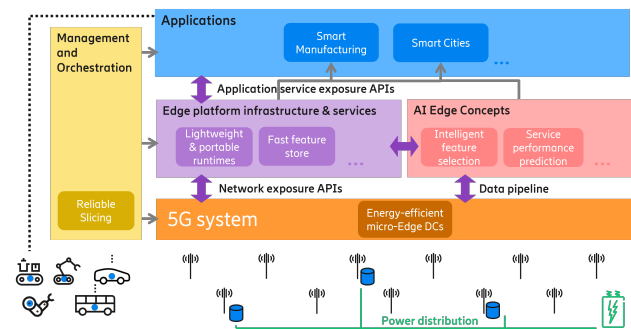


Figure 1: Functional structure of the ANIARA project.

Our studies include elaborating detailed use case descriptions for verticals using 5G and the edge cloud. ANIARA use cases cover two main aspects of factory operations, akin to [6]. These are: environmental monitoring and control of the factory floor such as air quality; temperature and power management and operations monitoring and control such as robot cell control, logistics and safety. These high-level use cases drive the system architecture requirements and serve as the basis for technology-specific use cases such as private 5G, edge micro datacenters and AI at the edge.

The remainder of this paper is organized to provide additional details of the components in Figure 1. Section 2 provides details about edge platform infrastructure and services that have been developing in ANIARA, including lightweight and portable execution environments and fast, dependable feature stores. Apart from edge infrastructure, developing edge AI enablers is also another important objective of the ANIARA. Section 3 describes AI/ML methods designed and developed within the project so far, including ML models for life-cycle management, intelligent feature selection and distributed learning in the edge considering privacy. ANIARA also addresses management and orchestration including both edge and cloud scenarios, to satisfy the specific needs of the

use cases, detailed in Section 4. Section 5 presents research on smart power for building a large-scale 5G edge system, shown in the lower portion of Figure 1. Section 6 provides a comprehensive related work. Finally, Section 7 gives an overview of works planned to be done in the future.

## 2 EDGE PLATFORM COMPONENTS

### 2.1 Programmable, light containerisation

Lightweight and portable execution environments have been identified as a crucial enabler for higher flexibility and dynamism of application deployment in a distributed network [24]. In ANIARA, we experiment with WebAssembly technologies to fill this gap. WebAssembly is an *open* binary instruction format for a stack-based virtual machine, designed to support existing programming languages in a web browser environment. In the browser, WebAssembly is generally faster than JavaScript due to its more compact format and manual memory management. The virtual machine is, however, not bound to browsers, but can be used standalone on various platforms. A deployment flow is where a high-level language such as C, C++, Rust, Python and so on can be compiled to WebAssembly (using clang) and executed on various hardware or software platforms. Webassembly allows edge applications to execute within a wide range of devices and operating systems. To prove the versatility and low footprint, we wrote an application in Rust and executed it on a WebAssembly runtime using a \$30 microcontroller. The same application was also written in Python that runs (on WebAssembly again) in an ASCII terminal.

### 2.2 Fast, dependable, feature store

A production-grade AI solution for edge computing, known as Feature stores, is a part of ANIARA. They act as a halfway house between data scientists and data engineers, enabling the same feature computation code to be used for model training and inference. Feature stores also act as a centralized repository of AI models<sup>1</sup> and have been adapted to perform in distributed scenarios for real-time data synchronization between geographically distributed Edge locations. The underlying database is called RonDB, evaluated using **LATS**, meaning low Latency, high Availability, high Throughput, scalable Storage principles to list RonDB’s performance:

- RonDB x3 times lower latency
- RonDB 1.1M reads/sec, REDIS 800k reads/sec
- RonDBs processed 200M reads in a 30 node cluster

## 3 AI FOR EDGE MANAGEMENT

We divide AI for edge management into three main research tracks: (i) ML model life-cycle management, (ii) Intelligent

feature selection, and (iii) Distributed learning in the edge. The three tracks have been studied within a use case for service performance prediction from edge statistics available to an operator with the objective to automate parts of the edge management process.

*ML model life-cycle management.* Ensuring high-performing ML models that are continuously updated and correctly deployed is critical for successful integration of ML in the edge. Transfer learning, which is one approach, allows for structured incorporation of previously acquired knowledge enabling timely and robust model adaptation, especially when data are scarce for reliable training of ML models. ANIARA studied concepts and methods for adapting and selecting ML models for mitigating ML-model performance degradation due to expected vertical and horizontal scaling in the edge infrastructure and 5G system, with specific focus on performance models for networked services, see [8, 13, 16],

*Intelligent feature selection.* A challenge for AI at the edge is related to network overhead with respect to measurements and monitoring, and feature selection for improved model performance [22]. A key enabler for ML models is timely access to reliable data, in terms of features, which require pervasive measurement points throughout the network. However, excessive monitoring is associated with network overhead. Using domain knowledge provides hints to find a balance between overhead reduction and identifying future ML requirements. ANIARA implemented a method of unsupervised feature selection that uses a structured approach in incorporation of the domain knowledge acquired from domain experts or previous learning experiences [21].

*Distributed learning in the edge.* Multi-domain service metric prediction is a key component in the ANIARA framework. It should enable privacy-preserving sharing of knowledge between operators, and low-overhead training of models within an operator in terms of data sharing. The approach requires in-network processing capabilities to enable federated learning. Therefore ANIARA is working on a multi-domain service metric prediction framework using federated learning, corresponding to a scenario where several services are managed by a number of operators in geographically distributed locations.

## 4 EDGE-CLOUD ORCHESTRATION

Slicing allows provisioning of multiple services over the same infrastructure, where virtual or physical resources are interconnected to form end-to-end logical networks, also known as slices [2]. Orchestrators run resource allocation algorithms to select the most suitable set of resources to satisfy the specific needs of the clients. In the edge cloud, compute

<sup>1</sup><https://www.featurestore.org>

resources are located close to the users, allowing provisioning of low latency services and enabling 5G Ultra-Reliable Low-Latency Communication slices. Allocating backup resources requires protecting the slices against link or node failures. Backup resources can be provided by means of a dedicated protection scheme, where resources are dedicated for each slice. Since backup resources are accessed only in case of failures, shared protection schemes can be developed, where backup resources are shared among different slices to decrease the overall amount of required resources.

Therefore, ANIARA has developed a heuristic-based shared protection to encourage sharing of backup connectivity and cloud resources. We evaluated this against a dedicated protection scheme using a Python simulator, published in [1]. Results show that the shared approach reduces the blocking probability by order of magnitude, and is especially beneficial when in-node processing resources are scarce.

## 5 EDGE-CLOUD POWER RESEARCH

Installing thousands of edge data centers, primarily in cities will require significant amounts of power, however many power grids are already utilized close to 100%. To maintain high availability, the edge-data centers need to be complemented with alternative power sources and pro-active power management systems. This requires tailored hardware solutions integrated with the power grid and on-site power generation. Supporting active load balancing by going off-grid for shorter periods of time. We are working on the design and implementation of a series of micro-edge-data center demonstrators for deployment at industrial sites. The first generation consists of a double rack configuration including, cooling, UPS-system with batteries, multiple power source inputs and IT-hardware.

To build out a large-scale 5G edge system, smart power utilization is required. One approach is to utilize the on-site UPS installation to go off-grid during peak power periods. The battery storage needs to be dimensioned to address this active usage. Incentivizing the active participation from the edge data centers in the load-balance activity is necessary. Moving away from fixed to dynamic prices will permits battery charging during off-peaks and discharging during higher-priced periods. Discharging implies less grid power.

The value for a power grid operator will be larger than that reflected by the customer energy price, if the data center is placed in a particularly energy-hungry section of the grid. Therefore, we have initiated a dialog with a major power grid operator.

## 6 RELATED WORK

The first incarnation of this paper appeared in [15] whilst other Edge roadmap papers include [5, 6]. A survey on transfer learning from methodological point of view is provided in [23], and [20] gives a recent survey on transfer learning for future wireless networks. A review on techniques for unsupervised feature selection is provided in [20]. A review on online feature selection techniques is given in [10]. The works in [7, 25] give review federated learning in mobile edge networks and edge computing, respectively. Our work builds upon these works, and extends the scope specifically towards edge clouds for telecom.

A seminal paper from four major browser companies who collaboratively designed WebAssembly can be found in [9]. A recent Ph.D describes an edge computing system which offloads computations from web-supported devices to edge server [11]. It exploits the portability of web apps when migrating the *execution state* of web apps. Edgedancer, a platform that offers infrastructure support for portable, provider-independent, and secure migration of edge services, it is a lightweight and generic execution environment by utilising WebAssembly, see [17]. A thorough survey on container technologies can be found in [3].

Resource allocation strategies for 5G networks and reliable services have been investigated recently. In particular, different techniques for backup protection of optical network resources, relying on both DP and SP schemes, have been presented in [14, 19]. Works propose efficient DP and SP algorithms for cloud and baseband resources in 5G access and metro networks, see [4, 12]. Considering connectivity and compute resources separately may lead to impractical solutions, especially when resources are scarce. Our work focuses on the dynamic slice provisioning where both type of resources are jointly allocated. A whitepaper presents an overview of the market and implementation trends, see [18].

## 7 FUTURE WORK

Going forward in the AI for edge field, we will study distributed learning under various sources of data and system heterogeneity. The objective is to tackle concerns with data privacy, resource heterogeneity among AI actors, challenges in re-usability of previously learned ML models, and difficulties in effective incorporation of domain knowledge. Experiments with WebAssembly based runtimes to implement *code-once, execute anywhere* approaches across the device-edge-cloud continuum are ongoing. The idea is to support offloading applications from the user equipment device to the edge node. Future work for the Feature store is a Kubernetes operator for RonDB and using it to store and serve our WASM containers. Improving RonDB and implementing an *evaluation store*, feature drift detection is planned. Before

a widespread edge data center can be used, we will need to work on power integration aspects that means deployment/installation of the physical hardware. Installation at hard-to-reach places, requiring easy assembly on-site and an autonomous operation with minimal on-site maintenance is ongoing (a demo was shown at the mid-term review, April 2022). We will also investigate the potentials and limitations of resource sharing in bare metal deployments of containers, and enhanced scaling strategies to improve utilization.

## ACKNOWLEDGEMENTS

We thank EU-CELTIC office under the project ID C2019/3-2. Also to the reviewers for feedback at mid-term review meeting, Berlin April 2022. We also acknowledge the Swedish, UK and Germany financing bodies: Vinnova AB "Automation of Network edge Infrastructure and Applications with aRTificial intelligence" 2020-00763, InnovateUK under the project ID 106197: ukANIARA, and the Bundesministerium für Bildung und Forschung under the name "AI-NET ANIARA 16KIF1274K". We thank Björn Skubic at Ericsson Research for his comments on an earlier version of this paper.

## REFERENCES

- [1] E. Amato, F. Tonini, C. Raffaelli, and P. Monti. 2021. A Resource Sharing Method for Reliable Slice as a Service Provisioning in 5G Metro Networks. In *2021 International Conference on Optical Network Design and Modeling (ONDM)*. 1–3.
- [2] Alcardo Alex Barakabitze, Arslan Ahmad, Rashid Mijumbi, and Andrew Hines. 2020. 5G network slicing using SDN and NFV: A survey of taxonomy, architectures and future challenges. *Computer Networks* 167 (2020), 106984. <https://doi.org/10.1016/j.comnet.2019.106984>
- [3] Emiliano Casalichio and Stefano Iannucci. 2020. The state-of-the-art in container technologies: Application, orchestration and security. *Concurrency and Computation: Practice and Experience* 32, 17 (2020), e5668. <https://doi.org/10.1002/cpe.5668>
- [4] H. D. Chantre and N. L. Saldanha da Fonseca. 2020. The Location Problem for the Provisioning of Protected Slices in NFV-Based MEC Infrastructure. *IEEE Journal on Selected Areas in Communications* 38, 7 (2020), 1505–1514. <https://doi.org/10.1109/JSAC.2020.2986869>
- [5] Aaron Ding, Ella Peltonen, Sasu Tarkoma, and Lars Wolf. 2021. Identifying Key Enablers in Edge Intelligence. *Dagstuhl Reports* 11, 7 (2021), 76–88. <https://doi.org/10.4230/DagRep.11.7.76>
- [6] Aaron Yi Ding, Ella Peltonen, Tobias Meuser, Atakan Aral, Christian Becker, Schahram Dustdar, Thomas Hiessl, Dieter Kranzlmüller, Madhusanka Liyanage, Setareh Magshudi, Nitinder Mohan, Jörg Ott, Jan S. Rellermeier, Stefan Schulte, Henning Schulzrinne, Gürkan Solmaz, Sasu Tarkoma, Blesson Varghese, and Lars C. Wolf. 2021. Roadmap for Edge AI: A Dagstuhl Perspective. *CoRR* abs/2112.00616 (2021). arXiv:2112.00616 <https://arxiv.org/abs/2112.00616>
- [7] Wei Yang Bryan Lim et al. 2020. Federated Learning in Mobile Edge Networks: A Comprehensive Survey. *IEEE Communications Surveys & Tutorials* 22 (2020), 2031–2063.
- [8] M. Ebrahimi F. G. Sanz and A. Johnsson. 2022. Exploring Approaches for Heterogeneous Transfer Learning in Edge Clouds. In *IEEE/IFIP Network Operations and Management Symposium (NOMS)*.
- [9] Andreas Haas, Andreas Rossberg, Derek L. Schuff, Ben L. Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon Zakai, and JF Bastien. 2017. Bringing the Web up to Speed with WebAssembly. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (Barcelona, Spain) (PLDI 2017)*. Association for Computing Machinery, New York, NY, USA, 185–200. <https://doi.org/10.1145/3062341.3062363>
- [10] Xuegang Hu, Peng Zhou, Peipei Li, Jing Wang, and Xindong Wu. 2016. A survey on online feature selection with streaming features. *Frontiers of Computer Science* 12 (2016), 479–493.
- [11] HyukJin Jeong. 2020. *Lightweight Offloading System For Edge Cloud Environments*. Ph.D. Dissertation. Department of Electrical and Computer Engineering The Graduate School Seoul National University.
- [12] Bahare M. Khorsandi, Federico Tonini, and Carla Raffaelli. 2019. Centralized vs. distributed algorithms for resilient 5G access networks. *Photonic Network Communications* 37, 3 (01 Jun 2019), 376–387.
- [13] Hannes Larsson, Jalil Taghia, Farnaz Moradi, and Andreas Johnsson. 2021. Source Selection in Transfer Learning for Improved Service Performance Predictions. In *2021 IFIP Networking Conference and Workshops*.
- [14] A. Marotta, D. Cassioli, M. Tornatore, Y. Hirota, Y. Awaji, and B. Mukherjee. 2020. Reliable Slicing with Isolation in Optical Metro-Aggregation Networks. In *2020 Optical Fiber Communications Conference and Exhibition (OFC)*. 1–3.
- [15] Ian Marsh, Nicolae Paladi, Henrik Abrahamsson, Jonas Gustafsson, Johan Sjöberg, Andreas Johnsson, Pontus Sköldström, Jim Dowling, Paolo Monti, Melina Vruna, and Mohsen Amiribesheli. [n.d.]. Evolving 5G: ANIARA, an edge-cloud perspective. *Proceedings of the 18th ACM International Conference on Computing Frontiers* ([n. d.]).
- [16] Farnaz Moradi, Rolf Stadler, and Andreas Johnsson. 2019. Performance Prediction in Dynamic Clouds using Transfer Learning. *2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)* (2019).
- [17] Manuel Nieke, Lennart Almstedt, and Rüdiger Kapitza. 2021. Edgedancer: Secure Mobile WebAssembly Services on the Edge. In *Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking (Online, United Kingdom) (EdgeSys '21)*. Association for Computing Machinery, New York, NY, USA, 13–18. <https://doi.org/10.1145/3434770.3459731>
- [18] The Linux Foundation Project. 2021. *State of the Edge 2021: A Market and Ecosystem Report for Edge Computing*. Whitepaper.
- [19] N. Shahriar, S. Taeb, S. R. Chowdhury, M. Zulfiqar, M. Tornatore, R. Boutaba, J. Mitra, and M. Hemmati. 2020. Reliable Slicing of 5G Transport Networks With Bandwidth Squeezing and Multi-Path Provisioning. *IEEE Transactions on Network and Service Management* 17, 3 (2020). <https://doi.org/10.1109/TNSM.2020.2992442>
- [20] Saúl Solorio-Fernández, Jesús Ariel Carrasco-Ochoa, and José Fco. Martínez-Trinidad. 2019. A review of unsupervised feature selection methods. *Artificial Intelligence Review* 53 (2019), 907–948.
- [21] Jalil Taghia, Farnaz Moradi, Hannes Larsson, Xiaoyu Lan, Masoumeh Ebrahimi, and Andreas Johnsson. 2022. Policy-Induced Unsupervised Feature Selection: A Networking Case Study. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*.
- [22] Xiaoxuan Wang, Forough Shahab Samani, Andreas Johnsson, and Rolf Stadler. 2021. Online Feature Selection for Low-overhead Learning in Networked Systems. In *2021 17th International Conference on Network and Service Management (CNSM)*. IEEE, 527–529.
- [23] Karl R. Weiss, Taghi M. Khoshgoftaar, and Dingding Wang. 2016. A survey of transfer learning. *Journal of Big Data* 3 (2016), 1–40.
- [24] G Wikström et al. 2022. 6G – Connecting a cyber-physical world: A research outlook towards 2030. *Ericsson, White paper* (Feb. 2022).
- [25] Qi Xia, Winson Ye, Zeyi Tao, Jindi Wu, and Qun Li. 2021. A survey of federated learning for edge computing: Research problems and solutions. *High-Confidence Computing* (2021).